The BGP TTL Security Hack (BTSH) <draft-gill-btsh-01.txt>

Vijay Gill, John Heasley, David Meyer

NANOG 27
Phoenix, Arizona

Agenda

- Problem Statement
- Can TTL be useful?
- Mechanisms
- A Few Usage Assumptions
- Limitations
- Questions/Comments
- http://www.maoz.com/~dmm/NANOG27/btsh

Problem Statement

- We're seeing a large number of DoS attacks against port 179. Why/How?
- Well, the TCP 4 tuple is easy enough to discover
- And the attack doesn't require sequence number. Why?
  - You don't care if the TCP segment valid. Just overload the RP. So...
- Conclusion: You don't have to own the attacked router to disable BGP processing

What can we (easily) do in the near term?

- i.e., is there anything (short of crypto techniques) we can use to mitigate these sorts of attacks?
- Well, TTL still (relatively) "hard" to spoof
- Can we somehow use that fact?

Is "TTL hard to spoof" useful?

- Well, since the vast majority of (eBGP) our peerings are between adjacent routers
- So why not set the TTL on BGP packets to 255 and then reject any BGP packets that come from configured peers which do NOT have a TTL in the range 255-254?
- That is, the receive TTL is expected to be within a small range of 1 or 2 (254-255). The actual value depends upon the (router) architecture, but it is expected that the receiver will verify the range

Is "TTL hard to spoof" useful?

- BTW, why not just set TTL=1 and let normal TTL=0 discard do its thing?

A Little More Mechanism

- Use a receive path ACL (e.g.) that only allows BGP packets to pass to the RP that have the correct <source,destination,TTL> tuple
  - i.e., TTL in the range 255-254 for directly connected peers
- If the TTL is not in the range 255-254
  - Punt the packet into low priority queue, log, or silently discard

A Few Usage Assumptions

- The vast majority of eBGP peerings are between adjacent routers
- It is common practice (or should be) for many service providers to ingress filter (deny) packets that have the provider's loopback addresses as the source IP address

A Few Usage Assumptions, cont.

- Use of BTSH is OPTIONAL, and can be configured on a per-peer/peer-group basis
- The router supports a method of classifying traffic destined for the route processor into interesting/control and not-control queues
- The peer routers both implement BTSH

Limitations

BTSH is a simple procedure that protects single hop BGP sessions, except in those cases where the directly connected peer has been compromised. While the method is less effective for multi-hop BGP sessions, it still closes the window on several forms of attack

And...

Protection of the BGP infrastructure beyond this method will likely require cryptographic machinery such as is envisioned by S-BGP (and/or other extensions)

Questions/Comments?

Thanks

Oh, and please attend the GROW BOF (Global Routing Operations Working Group) at the 56th IETF (Tuesday, March 18 at 1700-1800).